arXiv:1608.00346v1 [math.CO] 1 Aug 2016

WALKSAT STALLS WELL BELOW THE SATISFIABILITY THRESHOLD

AMIN COJA-OGHLAN, AMIR HAQSHENAS AND SAMUEL HETTERICH

ABSTRACT. Partly on the basis of heuristic arguments from physics it has been suggested that the performance of certain types of algorithms on random k-SAT formulas is linked to phase transitions that affect the geometry of the set of satisfying assignments. But beyond intuition there has been scant rigorous evidence that "practical" algorithms are affected by these phase transitions. In this paper we prove that `Walksat`, a popular randomised satisfiability algorithm, fails on random k-SAT formulas not very far above clause/variable density where the set of satisfying assignments shatters into tiny, well-separated clusters. Specifically, we prove `Walksat` is ineffective with high probability if $m/n > c2^k \ln^2 k/k$, where m is the number of clauses, n is the number of variables and $c > 0$ is an absolute constant. By comparison, `Walksat` is known to find satisfying assignments in linear time w.h.p. if $m/n < c'2^k/k$ for another constant $c' > 0$ [Coja-Oghlan and Frieze, SIAM J. Computing 2014].

Mathematics Subject Classification: 68Q87 (primary), 68W40 (secondary)

1. INTRODUCTION AND RESULTS

1.1. **Background and motivation.** For integers $k \geq 3$ and $n, m > 0$ let $\boldsymbol{\Phi} = \boldsymbol{\Phi}_k(n,m) = \boldsymbol{\Phi}_1 \wedge \ldots \wedge \boldsymbol{\Phi}_m$ be a random Boolean formula in conjunctive normal form with clauses $\boldsymbol{\Phi}_i = \boldsymbol{\Phi}_{i1} \vee \ldots \vee \boldsymbol{\Phi}_{ik}$ of length k over the Boolean variables $x_1, \ldots, x_n$ chosen uniformly at random from the set of all $(2n)^{km}$ possible such formulas. Since the very beginning research on the random k-SAT problem has been driven by two hypotheses. First, that for any $k \geq 3$ there is a certain critical clause-to-variable density $r_{k-\mathrm{SAT}} > 0$, the *k-SAT threshold*, where the probability that the random formula is satisfiable drops from almost 1 to nearly 0. Second, that random formulas with a density close to but below $r_{k-\mathrm{SAT}}$ are "computationally difficult" in some intuitive sense [10, 12, 33].

While over the past 20 years there has been tremendous progress on the first hypothesis [5, 6, 19, 17, 23], only few advances have been made on the second one. One exciting insight came from physics [31, 32]. Namely, according physics predictions, the geometry of the set of satisfying assignments $S(\boldsymbol{\Phi})$ undergoes a dramatic change well below the satisfiability threshold. Specifically, at a certain density m/n the set $S(\boldsymbol{\Phi})$ breaks up into tiny, well-separated "clusters" w.h.p. In fact, if we choose a satisfying assignment $\boldsymbol{\sigma}$ uniformly at random, then w.h.p. it will belong to a cluster with many "frozen variables", which take the same truth value in *all* the satisfying assignments in that cluster. Thus, the set $S(\boldsymbol{\Phi})$ is broadly resembles an error-correcting code, except that there is no simple underlying algebraic structure. In effect, if, say, a local search algorithm attempts to find a satisfying assignment, it would apparently have to have the foresight to steer into one cluster and get all its frozen variables right almost in one go. This appears impossible without a survey of the "global" dependencies amongst the variables.

Many of the physics predictions on the geometry of the set $S(\boldsymbol{\Phi})$, including "clustering" and "freezing", have by now become rigorous theorems. Moreover, the clause/variable density where clustering and freezing occur matches the density up to which algorithms are rigorously known to find satisfying assignments, at least asymptotically for large enough clause lengths k. To be precise, the k-SAT threshold is asymptotically equal to $m/n = 2^k \ln 2 - (1+\ln 2)/2 + o_k(1)$, where $o_k(1)$ hides a term that tends to 0 in the limit of large k [17, 19]. By comparison, algorithms are known to find satisfying assignments up to $m/n = (1+o_k(1))2^k \ln k/k$ [14]. Furthermore, for $m/n > (1+o_k(1))2^k \ln k/k$ clustering and freezing occur [3, 4, 34]. Thus, one might expect that random formulas turn "computationally difficult" for densities almost a factor of k below the k-SAT threshold. Yet despite the structural results and the compelling intuitive picture drafted by the physics work, it has emerged to be remarkably difficult to actually *prove* that these structural properties pose a barrier even for fairly simple satisfiability algorithms.

1.2. **The main result.** In this paper we provide such a proof for `Walksat`, one of the simplest non-trivial satisfiability algorithms. `Walksat` is a local search algorithm. It starts with a uniformly random assignment. So long as the current assignment fails to be satisfying, the algorithm chooses a random unsatisfied clause and flips the value assigned to a

⋆ The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 278857–PTCC.

Algorithm 1.1. `Walksat`(Φ,ω)
Input: A k-CNF Φ on V and an integer $\omega > 0$.
Output: A truth assignment.
1. Choose an initial assignment $\sigma^{[0]}$ uniformly at random.
2. For $i = 0,\ldots,\omega$ do
3. If $\sigma^{[i]}$ is a satisfying assignment output $\sigma^{[i]}$ and halt.
4. Choose $\Phi_i \in U_\Phi(\sigma^{[i]})$ and an integer from $1 \le j \le k$ uniformly at random.
5. Obtain $\sigma^{[i+1]}$ from $\sigma^{[i]}$ by flipping the value of the variable of the literal Φ_{ij}.
7. If $\sigma^{[\omega]}$ is a satisfying assignment output $\sigma^{[\omega]}$. Otherwise output 'failure'.

FIGURE 1. The `Walksat` algorithm.

random variable in that clause. That clause will thereby get satisfied, but other, previously satisfied clauses may become unsatisfied. If after a certain given number ω of iterations no satisfying assignment is found, `Walksat` gives up. Thus, the algorithm is one-sided: it may find a satisfying assignment but it cannot produce a certificate that a given formula is unsatisfiable. The pseudocode is shown in Figure 1; for a formula Φ with m clauses and $\sigma \in \Sigma$ we write $U_\Phi(\sigma)$ for the set of all indices $i \in [m]$ such that clause Φ_i is unsatisfied under σ and we let $\mathcal{U}_\Phi(\sigma) = |U_\Phi(\sigma)|$ be the number of unsatisfied clauses. `Walksat` is known to outperform exhaustive search by an exponential factor in the worst case and the procedure has been an ingredient for some of the best algorithms for the k-SAT problem [18, 24, 25, 27, 28, 36, 38].

For a given formula Φ and $\omega > 0$ we let $\mathrm{success}(\Phi,\omega)$ be the probability (over the random decisions of the algorithm only) that `Walksat`(Φ,ω) will find a satisfying assignment. Thus, $\mathrm{success}(\boldsymbol{\Phi},\omega)$ is a random variable that depends on the random formula $\boldsymbol{\Phi}$.

Theorem 1.2. *There is exists a constant $c > 0$ such that for all k and all $m/n \ge c2^k \ln^2 k/k$ w.h.p.*

$$\mathrm{success}(\boldsymbol{\Phi}, \lceil \exp(n/k^2) \rceil) \le \exp(-n/k^2).$$

The random formula $\boldsymbol{\Phi}$ is well-known to be unsatisfiable w.h.p. if $m/n > 2^k \ln 2$. Therefore, the condition $m/n > c2^k \ln^2 k/k$ in Theorem 1.2 implies a lower bound on the clause length k for which the statement is non-vacuous. We have not tried to optimise the constant c.

The density required by Theorem 1.2 exceeds the clustering/freezing threshold by a factor of $c \ln k$, but still the k-SAT threshold is almost a factor of k away. Moreover, the theorem shows that `Walksat` fails in a dramatic way: on typical random formula $\boldsymbol{\Phi}$ the success probability of `Walksat` is exponentially small, even if we run `Walksat` for an exponential number of rounds. In particular, even if we restart `Walksat` any polynomial number of times from a new starting point the cumulative success probability of all trials will remain exponentially small.

Why is it difficult to prove a result such as Theorem 1.2 given what we know about freezing/clustering? At the densities well below the k-SAT threshold like in Theorem 1.2 we know that a uniformly *random* satisfying truth assignment of the random formula $\boldsymbol{\Phi}$ will lie in a "frozen cluster" w.h.p. But there may very well exist unfrozen clusters; in fact, recent physics work suggests that there are exponentially many [9]. Hence, because `Walksat` just aims to find a single satisfying assignment rather than to sample one uniformly at random, the algorithm just needs to be lucky enough to find one weak, unfrozen spot, as it were. In other words, we have to rule out the possibility that the algorithm somehow manages to home in on those spots where the "barriers" of the set $S(\boldsymbol{\Phi})$ are easily overcome.

But establishing such a statement is well beyond the standard arguments for analysing algorithms on random structures. The main techniques such as the "method of differential equations" are suitable merely to trace algorithms for a small linear number of steps and run into severe difficulties if the algorithm ever backtracks. By construction, `Walksat` backtracks constantly (very likely many variables will likely be flipped more than once) and we actually need to follow the algorithm for an *exponential* number of steps. Hence, a different approach is needed. Section 2 provides a detailed outline of the proof of Theorem 1.2.

1.3. **Related work.** On the positive side, `Walksat` is known to find satisfying assignments for densities $m/n < 2^k/(25k)$ for large enough k in linear time [16]. Thus, the present paper matches the positve result up to a $O_k(\ln^2 k)$-factor. Physics arguments suggest that `Walksat` should actually be effective up to $m/n = (1+o_k(1))2^k/k$ [40], but not beyond. Positive results for `Walksat` for small k were obtained by Alekhnovich and Ben-Sasson [7]. Additionally, they obtained exponential lower bounds for `Walksat` in the planted 3-SAT problem for densities far above the

satisfiabilty threshold, where in the planted model a random 3-SAT formula is chosen conditioned on the existence of one solution [8].

Gamarnik and Sudan [22] obtained negative results for a class of algorithms that they call "sequential local algorithms" for the random k-NAESAT problem, a cousin of random k-SAT. Sequential local algorithms set the variables $x_1, \ldots, x_n$ of the random formula one by one in the natural order. They do not backtrack. The algorithm determines the value of variable x_i based on the depth-t neighborhood of x_i in the hypergraph respresenting the formula. To this end the algorithm takes into account the values assigned to those variables amongst $x_1, \ldots, x_{i-1}$ that occur in that part of the hypergraph. The class of sequential local algorithms encompasses truncated version of message passing algorithms such as Belief Propagation Guided Decimation and Survey Propagation Guided Decimation. 'Truncated' means that only a bounded number of parallel message updates are allowed; however, to reach an asymptotic fixed point of the messages it may be necessary to update for $\Theta(\ln n)$ rounds. The main result of [22] is that sequential local algorithms fail to find NAE-satisfying assignments for clause/variable densities above $C2^k \ln^2 k/k$ for a certain constant $C > 0$.

While `Walksat` is not a sequential local algorithm, we critically use one idea of the analysis from [22], called "overlap structures" in that paper. Specifically, Gamarnik and Sudan prove that for an appropriate integer l no l-tuple of NAE-satsifying assignments exist with pairwise distance about $n \ln(k)/k$ if the clause/variable densities is above $C2^k \ln^2 k/k$. However, a coupling argument shows that if a local sequential algorithm were likely to succeed, then there would have to be such an l-tuple of NAE-satisfying assignments with a non-vanishing probability. Actually the idea of overlap structures originates from the work of Rahman and Virag [37], who improved the density of an earlier negative result of Gamarnik and Sudan [21] for a more specialised class of algorithms for the independent set problem. The definition of "mists" in the present paper is directly inspired by overlap structures.

The first and the last author obtained negaitve results for message passing algorithms for random k-SAT that do not require bounds on the number of iterations [15, 26]. Specifically, [15] shows that a basic version of Belief Propagation Guided Decimation fails to find satisfying assignments for densities $m/n > C2^k/k$ for a certain constant $C > 0$. Moreover, [26] shows that a basic version of the conceptually more powerful Survey Propagation Guided Decimation algorithm fails if $m/n > (1 + o_k(1))2^k \ln k/k$.

Further negative results deal with DPLL-type algorithms. In particular, Achlioptas, Beame and Molloy [2] proved that certain types of DPLL-algorithms fail for densities $m/n > C2^k/k$. By comparison, unit clause propagation-type algorithms succeeds on random k-SAT formulas for $m/n < C'2^k/k$ [11, 13]. Finally, the best current algorithm for random k-SAT succeeds for $m/n \leq (1 + o_k(1))2^k \ln k/k$ but seems to fail beyond [14].

1.4. **Notation and preliminaries.** Throughout the paper we set $\rho = 2^{-k}m/n$ and $\kappa = \ln k/k$. We assume tacitly that k, n are sufficiently large for our various estimates to hold. Moreover, from here on out we always assume that $m/n > c2^k \ln^2 k/k$. (As mentioned above, the assumption that k is large is justified because we assume that $m/n > c2^k \ln^2 k/k$ and the random formula is unsatisfied w.h.p. if $m/n > 2^k \ln 2$.)

If l is a literal, then we write $|l|$ for the underlying variable. Thus, $|l| = x_i$ if $l = x_i$ or $l = \neg x_i$. Moreover, the Hamming distance of two truth assignments σ, τ is denoted by $\mathrm{dist}(\sigma, \tau)$. Additionally, for two truth assignments $\sigma, \tau : V \to \{0, 1\}$ we let

$$\Delta(\sigma, \tau) = \{x \in V : \sigma_1(x) \neq \tau(x)\} \tag{1.1}$$

be the set of variables where σ, τ differ; hence, $|\Delta(\sigma, \tau)| = \mathrm{dist}(\sigma, \tau)$. Further, for $\sigma \in \{0, 1\}^n$ and $r_1, r_2 \geq 0$ define

$$\mathcal{D}_\sigma(r_1, r_2) = \{\tau \in \Sigma : \lfloor r_1 \kappa n \rfloor \leq \mathrm{dist}(\sigma, \tau) \leq \lfloor r_2 \kappa n \rfloor\}. \tag{1.2}$$

Hence, $\mathcal{D}_\sigma(r_1, r_2)$ is a ring around σ with inner radius $r_1 \kappa n$ and outer radius $r_2 \kappa n$. Additionally, let $\mathcal{D}_\sigma(r) = \mathcal{D}_\sigma(r, r)$ be the set of assignments at distance exactly $r\kappa n$.

Recall that the *Kullback-Leibler divergence* of $p, q \in (0, 1)$ is defined as

$$D_{\mathrm{KL}}(q, p) = q \ln \frac{q}{p} + (1 - q) \ln \frac{1 - q}{1 - p}.$$

The following well-known lemma "Chernoff bound" states that the Kullback-Leibler divergence provides the rate function of the binomially distribution (e.g., [29, p. 21]).

Lemma 1.3. *Let $p, q \in (0,1)$ be distinct and let $X_n = \mathrm{Bin}(n,p)$. Then*

$$\lim_{n\to\infty} \frac{1}{n} \ln \mathbb{P}\left[X \le qn\right] = -D_{\mathrm{KL}}\left(q,p\right) \qquad \text{if } q < p,$$
$$\lim_{n\to\infty} \frac{1}{n} \ln \mathbb{P}\left[X \ge qn\right] = -D_{\mathrm{KL}}\left(q,p\right) \qquad \text{if } q > p.$$

We are going to need the following "random walk" version of Lemma 1.3.

Corollary 1.4. *Suppose that $(W_n)_{n\ge1}$ is a sequence of independent random variables such that $0 < \mathbb{P}[W_n = 1] = 1 - \mathbb{P}[W_n = -1] = p < 1/2$. Let $q > 0$. Then*

$$\lim_{n\to\infty} \frac{1}{n} \ln \mathbb{P}\left[\sum_{t=1}^{n} W_n \ge qn\right] = -D_{\mathrm{KL}}\left((1+q)/2, p\right).$$

Proof. Let $X_t = (1+W_t)/2$ for all $t \ge 1$. Then $S_n = \sum_{t=1}^n X_t$ is a binomial random variable with parameters n and p and $\sum_{t=1}^n W_t = 2(\sum_{t=1}^n X_t) - n$. Hence, $\sum_{t=1}^n W_t \ge qn$ iff $\sum_{t=1}^n X_t \ge n(1+q)/2$ and the assertion follows from Lemma 1.3. □

2. Outline

The classical worst-case analysis of `Walksat` goes as follows. Suppose that Φ is a satisfiable k-SAT formula on n variables and fix a satisfying assignment τ. At any step the algorithm flips a randomly chosen variable in an unsatisfied clause. Because τ must satisfy that clause, there is at least a $1/k$ chance that the algorithm moves toward τ. Hence, in the case $k = 2$ the distance evolves at least as good as in an unbiased random walk, and thus we expect to reach τ or another satisfying assignment in $O(n^2)$ steps [35]. By contrast, for $k \ge 3$ the corresponding random walk has a drift away from τ and the probability of reaching τ in polyonmial time from a random starting point is exponentially small. Yet calculating the probability of starting at distance a bit less than $n/2$ from τ and then dashing towards it reveals that `Walksat` beats the naive 2^n exhaustive search algorithm [38].

Of course, on a random formula this analysis is far from tight. For example, for m/n below the satisfiability threshold the number $|S(\mathbf{\Phi})|$ of satisfying assignments is typically exponential in n. In fact, w.h.p. we have $\ln|S(\mathbf{\Phi})| = n\ln 2 + \frac{m}{n}\ln(1-(1+o_k(1))2^{-k})$ [4]. Hence, if $m/n = O_k(2^k \ln^2 k/k)$, then w.h.p. the number of satisfying assignments is as large as

$$|S(\mathbf{\Phi})| = 2^{n(1-O_k(\ln^2 k/k))}.$$

This observation obliterates some obvious proof ideas, such as combining the random walk argument from the previous paragraph with some sort of a union bound on the number of satisfying assignments; there is just too many of them.[1]

Another type of approach that seems doomed is meticilously tracing every step of the `Walksat` algorithm. This is basically what the proof of the positive `Walksat` result from [16] does. Such analyses typically depend on the principle of deferred decisions, i.e., the idea that the parts of the formula that the algorithm has not inspected yet are "random", subject to some relatively weak conditioning. This kind of approach can follow an algorithm for a small linear number of steps. But here we are trying to prove a statement about an exponential number of iterations. By that time the algorithm will likely have visited every clause of the formula several times over and thus there is "no randomness left". Hence, we need a different approach.

Our strategy is to split the analysis in two parts. First, we are going to formulate a few *quasirandom properties*. We will show that `Walksat` is exponential on *any* given formula that has these properties. Second, we will prove that the random formula has these quasirandom properties w.h.p. A similar type of argument was used, e.g., in prior work on message passing algorithms [15, 26].

The key is to come up with the right quasirandom properties. To this end, we need to develop an intuition as to what `Walksat` actually does on a random input $\mathbf{\Phi}$. Because `Walksat` starts from a random assignment, initially there will be about $2^{-k}m = \rho n$ unsatisfied clauses. In fact, we can establish a stronger, more geometric statement. Let $T(\Phi)$ be the set of all truth assignments $\tau \in \{0,1\}^n$ such that $\mathcal{U}_{\mathbf{\Phi}}(\tau) \le n\rho/10$ (i.e., the number of violated clauses is a tenth of what we expect in a random assignment). Set $\kappa = \ln k/k$. Then a union bound shows that the initial assignment $\sigma^{[0]}$ will most likely be at distance at least $10\kappa n$ from all $\tau \in T(\mathbf{\Phi}) \supset S(\mathbf{\Phi})$.

[1] The second author's master thesis contained an argument based on combining the random walk analysis with a union bound. However, that argument requires that $m/n = (1+o_k(1))2^k \ln 2$, a much stronger assumption than that of Theorem 1.2.

The second observation is that `Walksat` will likely have a hard time entering the set $T(\mathbf{\Phi})$. Intuitively, for $m/n > (1+o_k(1))2^k \ln k/k$ it is not just the set $S(\mathbf{\Phi})$ that shatters into tiny well-separated clusters, but even the set $T(\mathbf{\Phi})$ has this property. Moreover, the no man's land between different clusters provides no clues that nudge `Walksat` towards any one of them. In fact, there is a repulsion effect. To be precise, consider a "target assignment" $\tau \in T(\mathbf{\Phi})$ and suppose that $\sigma \in \{0,1\}^n \setminus T(\mathbf{\Phi})$ has distance at most $100\kappa n$ from τ. Because $\sigma \notin T(\mathbf{\Phi})$, the assignment leaves at least $n\rho/10$ clauses unsatisfied. Let us pretend that these unsatisfied clauses are random. Then if we pick a variable in an unsatisfied clause randomly, the probability of hitting a variable in $\Delta(\sigma,\tau)$ is as small as $100\kappa < 0.1$ (for large enough k). Hence, there is a 90% chance that `Walksat` will move away from τ, deeper into no man's land. Thus, to reach a satisfying assignment or, in fact any assignment in $T(\mathbf{\Phi})$ `Walksat` would have to beat the odds and overcome a substantial negative drift, which is exponentially unlikely.

However, there is one point that we missed. Although the probability of walking towards one satisfying assignment at distance at most $100\kappa n$ from the present assignment may be small, the total number of satisfying assignments is enormous and `Walksat` just has to find any one of them. In other words, at any step `Walksat` may be taking part in an exponential number of "lotteries". While any one of them may be rigged against the algorithm, the sheer number of simultaneous lotteries may yet give the algorithm a chance to succeed in polynomial time.

To rule this possibility out we introduce the concept of a *mist*, which is an adaptation of the "overlap structures" from [22]. More precisely, we will argue that we do not need to track the distance between `Walksat`'s current assignment and the entire set $T(\mathbf{\Phi})$ but merely the distance to a much smaller set $\mathcal{M}$ of assignments. This subset is "sparse" in the sense that for any truth assignment σ the number of assignments in $\mathcal{M}$ at distance at most $10\kappa n$ from σ is bounded by k rather than exponential in n. We will use this fact to argue that at any time the algorithm only takes part in at most k lotteries rather than an exponential number. This will enable us to prove that reaching $T(\mathbf{\Phi})$ will most likely take an exponential amount of time.

Formally, let Φ be a k-CNF on the variable set $x_1,\ldots,x_n$.

A *mist* of Φ is a set $\mathcal{M} \subset T(\Phi)$ of assignments with the following two properties.

MI1: the assignments in $\mathcal{M}$ have pairwise distance at least $2\kappa n$.
MI2: for each $\sigma \in T(\Phi)$ there exists $\mu \in \mathcal{M}$ such that $\mathrm{dist}(\mu,\sigma) \le 2\kappa n$.

Thus, the points of the mist are spread out but there is one near every assignment in $T(\Phi)$. Let

$$\mathcal{D}(\Phi,\mathcal{M}) = \bigcup_{\sigma\in\mathcal{M}} \mathcal{D}_\sigma(0,10)$$

be the set of all assignments at distance at most $10\kappa n$ from $\mathcal{M}$. Moreover, for a truth assignment σ and a set $W \subset \{x_1,\ldots,x_n\}$ let

$$X_\Phi(W,\sigma) = \sum_{i\in U_\Phi(\sigma)} \sum_{j\in[k]} \mathbf{1}\{|\Phi_{ij}| \in W\} \tag{2.1}$$

be the number of occurrences of variables from W in the unsatisfied clauses $U_\Phi(\sigma)$. Further, call Φ *quasirandom* if there is a mist $\mathcal{M}$ such that the following three statements hold.

Q1: we have $|\mathcal{D}(\Phi,\mathcal{M})| \le 2^n \exp(-2n/k^2)$.
Q2: for any $\tau \in \{0,1\}^n$ we have $|\mathcal{M} \cap \mathcal{D}_\tau(0,10)| \le k$.
Q3: for every $\mu \in \mathcal{M}$ in the mist and each $\sigma \in \mathcal{D}_\mu(0,100) \setminus T(\Phi)$ we have $X_\Phi(\Delta(\mu,\sigma)) \le k\mathcal{U}_\Phi(\sigma)/10$.

Thus, the set $\mathcal{D}(\Phi,\mathcal{M})$ is small and thus it is exponentially unlikely for the initial random $\sigma^{[0]}$ to belong to this set. Moreover, there are no more than k elements of the mist $\mathcal{M}$ in the vicinity of any one assignment τ. Finally, **Q3** says that if $\tau \notin T(\Phi)$ is an assignment with many unsatisfied clauses at distance no more than $100\kappa n$ from $\mu \in \mathcal{M}$, then the probability that `Walksat` takes a step from τ towards μ does not exceed 10%. Indeed, $X_\Phi(\Delta(\mu,\tau))$ is the number of flips that take `Walksat` closer to μ, and $k\mathcal{U}_\Phi(\tau)$ is the total number of possible flips.

Now, proving Theorem 1.2 comes down to establishing the following two statements.

Proposition 2.1. *If Φ is quasirandom, then* $\mathrm{success}(\Phi, \lceil \exp(n/k^2)\rceil)] \le \exp(-n/k^2)$.

Proposition 2.2. *If $m/n \ge 195\cdot 2^k \ln^2 k/k$, then $\mathbf{\Phi}$ is quasirandom w.h.p.*

We prove Proposition 2.1 in Section 3 and Proposition 2.2 in Section 4. Theorem 1.2 is immediate from Propositions 2.1 and 2.2.

3. Proof of Proposition 2.1

Suppose that $\Phi = \Phi_1 \wedge \cdots \wedge \Phi_m$ is a quasirandom k-CNF on the variables $x_1, \ldots, x_n$. Let $\mathcal{M}$ be a mist such that **Q1–Q3** hold and set $\omega = \lceil \exp(n/k^2) \rceil$. Condition **Q1** provides that the event $\mathcal{A} = \{\sigma^{[0]} \notin \mathcal{D}(\Phi)\}$ has probability

$$\mathbb{P}[\mathcal{A}] \geq 1 - \exp(-2n/k^2). \tag{3.1}$$

In the following we may therefore condition on $\mathcal{A}$.

The key object of the proof is the following family of events: for $\mu \in \mathcal{M}$ and $1 \leq t_1 < t_2 \leq \omega$ let

$$H_\mu(t_1, t_2) = \left\{ \mathrm{dist}(\sigma^{[t_1]}, \mu) = \lfloor 10\kappa n \rfloor, \mathrm{dist}(\sigma^{[t_2]}, \mu) = \lfloor 5\kappa n \rfloor, \forall t_1 \leq t \leq t_2 : \sigma^{[t]} \in \mathcal{D}_\mu(5, 10) \setminus T(\Phi) \right\}. \tag{3.2}$$

In words, $H_\mu(t_1, t_2)$ is the event that at time t_1 `Walksat` stands at distance precisely $\lfloor 10\kappa n \rfloor$ from μ, that the algorithm advances to distance $\lfloor 5\kappa n \rfloor$ at time t_2 while not treading closer to μ but staying in $\mathcal{D}_\mu(5, 10)$ at any intermediate step, and that `Walksat` does not hit $T(\Phi)$ at any intermediate step. Let

$$\mathcal{H} = \bigcup_{\mu \in \mathcal{M}, 0 \leq t_1 < t_2 \leq \omega} H_\mu(t_1, t_2).$$

Fact 3.1. *We have* $\mathbb{P}\left[\exists t \leq \omega : \sigma^{[t]} \in S(\Phi) | \mathcal{A}\right] \leq \mathbb{P}[\mathcal{H} | \mathcal{A}]$.

Proof. Recall that $S(\Phi) \subset T(\Phi)$. Suppose that $\sigma^{[t]} \in S(\Phi)$ for some $t \leq \omega$; then the algorithm halts at time t. Let $t_0 < t$ be minimum such that $\sigma^{[t_0]} \in T(\Phi)$. Then there exists $\mu \in \mathcal{M}$ such that $\mathrm{dist}(\mu, \sigma^{[t_0]}) < 2\kappa n$. Further, given $\mathcal{A}$ we have $\mathrm{dist}(\sigma^{[0]}, \mu) > 10\kappa n$. Hence, for some $0 < t_1 < t_0$ the event $\mathrm{dist}(\sigma^{[t_1]}, \mu) \leq 10\kappa n$ occurs for the first time. Moreover, there exists a minimum t_2 such that $t_1 < t_2 < t_0$ and $\mathrm{dist}(\sigma^{[t_2]}, \mu) \leq 5\kappa n$. Since `Walksat` moves Hamming distance one in each step, $H_\mu(t_1, t_2)$ occurs. □

To show that $\mathcal{H}$ is exponentially unlikely we are first going to estimate the probability of a single event $H_\mu(t_1, t_2)$.

Lemma 3.2. *Let* $\tau_1 \notin T(\Phi)$ *and* $\mu \in \mathcal{M}$ *be such that dist*$(\tau_1, \mu) = \lfloor 10\kappa n \rfloor$. *Then*

$$\mathbb{P}\left[H_\mu(t_1, t_2) | \mathcal{A}, \sigma^{[t_1]} = \tau_1\right] \leq \exp(-\kappa n/2) \quad \textit{for all } 1 \leq t_1 \leq t_2 \leq \omega.$$

Proof. For an index $t_1 < t \leq t_2$ define

$$Y_{t+1} = \mathrm{dist}(\sigma^{[t+1]}, \mu) - \mathrm{dist}(\sigma^{[t]}, \mu) + 2 \cdot \mathbf{1}\{\sigma^{[t]} \notin \mathcal{D}_\mu(5, 10) \setminus T(\Phi)\}. \tag{3.3}$$

If the event $H_\mu(t_1, t_2)$ occurs, then $\sigma^{[t]} \in \mathcal{D}_\mu(5, 10) \setminus T(\Phi)$ for all $t_1 \leq t \leq t_2$ and $\sum_{t_1 \leq t < t_2} Y_{t+1} \leq 1 - 5\kappa n$. Moreover, we claim that

$$\mathbb{E}[Y_{t+1} - Y_t | \sigma^{[t]} \notin T(\Phi)] \geq 4/5. \tag{3.4}$$

Indeed, at time $t+1$ `Walksat` chooses an unsatisfied clause and then a variable from that clause uniformly at random. If $Y_{t+1} < Y_t$, then the chosen variable is from the set $\Delta(\mu, \sigma^{[t]})$ of variables where $\sigma^{[t]}$ and μ differ. By (2.1) the probability of this event equals $X_\Phi(W, \sigma^{[t]})/k\mathcal{U}_\Phi(\sigma^{[t]})$. Hence, **Q3** shows that the probability that $\mathrm{dist}(\sigma^{[t+1]}, \mu) < \mathrm{dist}(\sigma^{[t]}, \mu)$ is bounded by 0.1, unless $\sigma^{[t]} \notin \mathcal{D}_\mu(5, 10) \setminus T(\Phi)$. Consequently, (3.4) follows from the definition (3.3).

If we let $(W_t)_{t \geq 1}$ be a sequence of independent ± 1-random variables such that $\mathbb{P}[W_t = -1] = 0.1$ and $\mathbb{P}[W_t = 1] = 0.9$, then (3.4) implies

$$\mathbb{P}\left[H_\mu(t_1, t_2) | \mathcal{A}, \sigma^{[t_1]} = \tau_1\right] \leq \mathbb{P}\left[\sum_{t_1 \leq t < t_2} Y_{t+1} \leq 1 - 5n \ln k/k\right] \leq \mathbb{P}\left[\sum_{t_1 \leq t < t_2} W_t \leq 1 - 5n \ln k/k\right].$$

Thus, the assertion follows from Corollary 1.4 and the fact that $H_\mu(t_1, t_2)$ can occur only if $t_2 - t_1 \geq 5\kappa n$, because `Walksat` moves Hamming distance one in each step. □

Proof of Proposition 2.1. By Lemma 3.2 each of the events contributing to $\mathcal{H}$ occurs only with probability at most $\exp(-\kappa n/2)$ given $\mathcal{A}$. But since the number of assignments in the mist $\mathcal{M}$ and hence the number of individual events $H_\mu(t_1, t_2)$ may be much larger than $\exp(n\kappa/2)$, a simple union bound on $\mu \in \mathcal{M}$ won't do. Indeed, the real problem

here is the size of the mist and not the number of possible choices of t_1, t_2, because $t_1, t_2 \leq \omega$ and ω is (exponential but) relatively small. In other words, we do not give away too much by writing

$$\begin{aligned}\mathbb{P}\left[\mathcal{H}|\mathcal{A}\right] &\leq \sum_{0\leq t_1<t_2\leq\omega} \mathbb{P}\left[\bigcup_{\mu\in\mathcal{M}} H_\mu(t_1,t_2)\middle|\mathcal{A}\right] \\ &= \sum_{0\leq t_1<t_2\leq\omega}\sum_{\sigma\in\Sigma} \mathbb{P}\left[\bigcup_{\mu\in\mathcal{M}} H_\mu(t_1,t_2)|\mathcal{A},\sigma^{[t_1]}=\sigma\right]\mathbb{P}\left[\sigma^{[t_1]}=\sigma|\mathcal{A}\right] \\ &\leq \sum_{0\leq t_1<t_2\leq\omega} \max_{\sigma\in\Sigma}\mathbb{P}\left[\bigcup_{\mu\in\mathcal{M}} H_\mu(t_1,t_2)|\mathcal{A},\sigma^{[t_1]}=\sigma\right] \\ &\leq \sum_{0\leq t_1<t_2\leq\omega} \max_{\sigma\in\Sigma}\sum_{\mu\in\mathcal{M}}\mathbb{P}\left[H_\mu(t_1,t_2)|\mathcal{A},\sigma^{[t_1]}=\sigma\right]. \end{aligned} \tag{3.5}$$

To bound the last term, we recall from (3.2) that $\mathbb{P}\left[H_\mu(t_1,t_2)|\mathcal{A},\sigma^{[t_1]}=\sigma\right]=0$ unless $\mathrm{dist}(\mu,\sigma)=\lfloor 10\kappa n\rfloor$. Hence, **Q2** implies that for any $\sigma\in\Sigma$ the sum on μ in (3.5) has at most k non-zero summands. Therefore, Lemma 3.2 gives

$$\max_{\sigma\in\Sigma}\sum_{\mu\in\mathcal{M}}\mathbb{P}\left[H_\mu(t_1,t_2)|\mathcal{A},\sigma^{[t_1]}=\sigma\right]\leq k\exp(-n\kappa/2). \tag{3.6}$$

Plugging (3.5) into (3.6) and recalling the choice of ω, we get

$$\mathbb{P}\left[\mathcal{H}|\mathcal{A}\right]\leq\omega^2 k\exp(-\kappa n/2)\leq\exp(-n/k^2), \tag{3.7}$$

with room to spare. Finally, the assertion follows from (3.1), Fact 3.1 and (3.7). □

4. Proof of Proposition 2.2

We begin with the following standard 'first moment' bound.

Lemma 4.1. *We have* $\mathbb{E}\left|T(\boldsymbol{\Phi})\right|\leq 2^n\exp\left(-\rho n/2\right)$.

Proof. For any fixed assignment $\sigma\in\{0,1\}^n$ the number $\mathcal{U}_{\boldsymbol{\Phi}}(\sigma)$ of unsatisfied clauses has distribution $\mathrm{Bin}(m,2^{-k})$. Therefore, by Lemma 1.3 and our assumption on m/n,

$$\mathbb{P}\left[\sigma\in T(\boldsymbol{\Phi})\right]=\exp(-mD_{\mathrm{KL}}\left(0.1\cdot 2^{-k},2^{-k}\right)+o(n))\leq\exp(-\rho n/2).$$

Thus, the assertion follows from the linearity of expectation. □

To proceed, we construct a mist $\mathcal{M}$ of the random formula $\boldsymbol{\Phi}$ by means of the following iterative procedure.

(1) Initially let $\mathcal{M}=\emptyset$.
(2) While $T(\boldsymbol{\Phi})\setminus\bigcup_{\mu\in\mathcal{M}}\mathcal{D}_\mu(0,2)\neq\emptyset$, add an arbitary element of this set to $\mathcal{M}$.

Let us fix any possible outcome $\mathcal{M}$ of the above process. Of course, $\mathcal{M}$ depends on $\boldsymbol{\Phi}$ but we do not make this explicit to unclutter the notation. We now simply verify the conditions **Q1–Q3** one by one.

Lemma 4.2. ***Q1*** *holds with probability* $1-\exp(-\Omega(n))$

Proof. We start with a naive bound on the number of assignments in $\mathcal{D}_\sigma(0,10)$ centered at an arbitrary $\sigma\in\Sigma$. Stirling's formula shows that for any fixed assignment $\sigma\in\{0,1\}^n$,

$$|\mathcal{D}_\sigma(0,10)|\leq\sum_{j\leq 10\kappa n}\binom{n}{j}\leq n\exp(10n\ln^2 k/k).$$

Hence, the construction of $\mathcal{M}$ ensures that $|\mathcal{D}(\boldsymbol{\Phi},\mathcal{M})|\leq|T(\boldsymbol{\Phi})|\cdot n\exp(10n\ln^2 k/k)$. Thus,

$$\mathbb{E}\left[|\mathcal{D}(\boldsymbol{\Phi})|\right]\leq\mathbb{E}\left[|T(\boldsymbol{\Phi})|\right]\cdot n\exp(10n\ln^2(k)/k).$$

Consequently, the assertion follows from Lemma 4.1 and our assumption on ρ. □

For an assignment $\sigma\in\Sigma$ let $\mathcal{C}(\sigma)$ be the set of all possible unsatisfied clauses under σ on the variable set $x_1,\ldots,x_n$. Then $|\mathcal{C}(\sigma)|=n^k$ for all $\sigma\in\Sigma$.

The following Lemma proving that with high probability **Q2** holds in $\boldsymbol{\Phi}$ is similar to the statement in [22] that certain "overlap structures" do not exist (where an "overlap structure" is an l-tuple of NAE-satisfying assignments with pairwise distance $\sim\kappa n$ for an appropriate integer l.) This concept is an adaption of a bound on intersection

densities for tuples of independent sets in sparse d-regular graphs from [37]. There it is shown that no tuple of large local independent sets intersecting each other in a certain way exists in a d-regular graph w.h.p. We are going to prove a similar statement, namely that no m-tuple of assignments with a small number of unsatisfied clauses that have pairwise distance $\sim \kappa$ and are all contained in $\mathcal{D}_\tau(0,10)$ for some $\tau \in \Sigma$ exist. Following [22] we also use an inclusion/exclusion estimate, while here of course we are not focussing on satisfying assignments but on assignments with a relatively small number of unsatisfied clauses.

Lemma 4.3. ***Q2*** *holds w.h.p.*

Proof. We prove the statement by way of a slightly different random formula model $\mathbf{\Phi}'$. In $\mathbf{\Phi}'$ each of the $(2n)^k$ possible clauses is included with probability $q = m/(2n)^k$ independently in a random order. A standard argument shows that this model is essentially equivalent to $\mathbf{\Phi}$. To be precise, we claim that for any event $\mathcal{E}$ we have

$$\mathbb{P}\left[\mathbf{\Phi} \in \mathcal{E}\right] \leq O(\sqrt{n})\,\mathbb{P}\left[\mathbf{\Phi}' \in \mathcal{E}\right] + o(1). \tag{4.1}$$

To see this, let $\mathcal{G}$ be the event that $\mathbf{\Phi}$ does not contain the same k-clause twice, i.e., $\mathbf{\Phi}_i \neq \mathbf{\Phi}_j$ for all $1 \leq i < j \leq m$. A simple union bound shows that $\mathbb{P}\left[\mathbf{\Phi} \in \mathcal{G}\right] = 1 - O(1/n)$. Moreover, let $\boldsymbol{m}'$ be the total number of clauses of $\mathbf{\Phi}'$. The $\boldsymbol{m}'$ is a binomial variable with mean m and Stirling's formula shows that $\mathbb{P}\left[\boldsymbol{m}' = m\right] = \Theta(n^{-1/2})$. Thus, (4.1) follows from the observation that the distribution of $\mathbf{\Phi}'$ given $\boldsymbol{m}' = m$ coincides with the distribution of $\mathbf{\Phi}$ given $\mathcal{G}$.

Hence, we are going to work with the model $\mathbf{\Phi}'$. Let $\mathcal{M}'$ be the mist constructed for $\mathbf{\Phi}'$ by means of our above procedure. Moreover, for $\tau \in \{0,1\}^n$ let $P(\tau)$ be the set of all k-tuples $(\sigma_i)_{i\in[k]}$ with the following two properties.

P1: $\sigma_i \in \mathcal{D}_\tau(0,10)$ for all $i \in [k]$ and
P2: $\mathrm{dist}(\sigma_i, \sigma_j) \geq 2n\kappa$ for all $i \neq j$.

Then

$$|P(\tau)| \leq n \cdot \binom{n}{n10\ln(k)/k}^k \leq n \cdot \left(\frac{\mathrm{e}k}{10\ln k}\right)^{10n\ln k} \leq \exp\left(10n\ln^2 k\right). \tag{4.2}$$

Further, if $\sigma_1, \sigma_2 \in \{0,1\}^n$ are assignments such that $\mathrm{dist}(\sigma_1,\sigma_2) \geq 2\kappa n$, then the number of possible unsatisfied clauses under both σ_1 and σ_2 satisfies

$$|C(\sigma_1) \cap C(\sigma_2)| = (n - \mathrm{dist}(\sigma_1,\sigma_2))^k \leq ((1 - 2\ln(k)/k)n)^k \leq k^{-2}n^k; \tag{4.3}$$

this is because a clause that is unsatisfied under both σ_1, σ_2 must not contain any literals on which the two assignments differ. We are going to upper bound the probability that for $(\sigma_i)_{i\in[k]} \in P(\tau)$ assignment σ_i renders at most $\rho n/10$ clauses of $\mathbf{\Phi}'$ unsatisfied given that all of $\sigma_1, \ldots, \sigma_{i-1}$ do so any $i = 1, \ldots, k$. The probability that this event occurs is upper bounded by the probability that $\mathbf{\Phi}'$ contains at most $\rho n/10$ clauses from the set

$$C(\sigma_i|\sigma_1,\ldots,\sigma_{i-1}) = C(\sigma_i) \setminus \bigcup_{j=1}^{i-1} C(\sigma_j). \tag{4.4}$$

The estimate (4.3) and inclusion/exclusion yield

$$|C(\sigma_i|\sigma_1,\ldots,\sigma_{i-1})| \geq n^k(1 - (i-1)k^{-2}).$$

Hence, if we let $Z_i = \mathrm{Bin}(\lfloor n^k(1-(i-1)k^{-2})\rfloor, q)$, then

$$\mathbb{P}\left[\sigma_i \in T(\boldsymbol{\Phi}')|\sigma_1,\ldots,\sigma_{i-1} \in T(\boldsymbol{\Phi}')\right] \leq \mathbb{P}\left[Z_i \leq \rho n/10\right] \tag{4.5}$$

(this step required that the clauses of $\boldsymbol{\Phi}'$ appear independently). By the Chernoff bound, for $i \leq k$ we have

$$\mathbb{P}\left[Z_i \leq \rho n/10\right] \leq \exp\left(-\rho n/15\right) \tag{4.6}$$

Consequently, **P2**, (4.5) and (4.6) yield for any $(\sigma_i)_{i\in[k]} \in P(\tau)$,

$$\mathbb{P}\left[\sigma_1,\ldots,\sigma_k \in T(\boldsymbol{\Phi})\right] = \prod_{i=1}^k \mathbb{P}\left[\sigma_i \in T(\boldsymbol{\Phi})|\sigma_j \in T(\boldsymbol{\Phi}) \text{ for all } j < i\right] \leq \exp\left(-k\rho n/15\right). \tag{4.7}$$

Further, let $Q(\boldsymbol{\Phi}', \tau)$ be the set of all k-tuples $(\sigma_i)_{i\in[k]} \in P(\tau)$ such that $\sigma_1,\ldots,\sigma_k \in T(\boldsymbol{\Phi}')$. Then (4.2) and (4.7) imply

$$\mathbb{E}\left[Q(\boldsymbol{\Phi}', \tau)\right] \leq \exp\left[n\left(10\ln^2(k) - k\rho/15\right)\right]. \tag{4.8}$$

Summing (4.8) on $\tau \in \{0,1\}^n$ and using $\rho \geq 195\ln^2(k)/k$, we get

$$\sum_{\tau\in\{0,1\}^n} \mathbb{E}\left|Q(\boldsymbol{\Phi}',\tau)\right| \leq \exp\left[n\left((2+10)\ln^2(k) - 13\ln^2(k)\right)\right] = \exp(-\Omega(n)). \tag{4.9}$$

Finally, assume that $\boldsymbol{\Phi}'$ violates **Q2**. Then there is $\tau \in \{0,1\}^n$ such that $Q(\boldsymbol{\Phi}',\tau) \neq \emptyset$, because our construction of $\mathcal{M}'$ ensures that $\mathcal{M}' \subset T(\boldsymbol{\Phi}')$ and that the pairwise distance of assignments in $\mathcal{M}$ is at least $2n\kappa$. Consequently, (4.9) shows together with Markov's inequality that $\boldsymbol{\Phi}'$ violates **Q2** with probability at most $\exp(-\Omega(n))$. Thus, the assertion follows by transferring this result to $\boldsymbol{\Phi}$ via (4.1). □

Lemma 4.4. *$\boldsymbol{\Phi}$ satisfies **Q3** w.h.p.*

Proof. Let $\mathcal{P} = \mathcal{P}_{\boldsymbol{\Phi}}$ be the number of pairs $(\sigma,\tau) \in \{0,1\}^n \times (\mathcal{D}_\sigma(0,100) \setminus T(\boldsymbol{\Phi}))$ such that $X_{\boldsymbol{\Phi}}(\Delta(\sigma,\tau)) > k\mathcal{U}_{\boldsymbol{\Phi}}(\tau)/10$. To estimate $\mathcal{P}$ fix a pair (σ,τ) and let $\mathcal{P}_{\boldsymbol{\Phi}}(\sigma,\tau)$ be the event that $X_{\boldsymbol{\Phi}}(\Delta(\sigma,\tau)) > k\mathcal{U}_{\boldsymbol{\Phi}}(\tau)/10$. If $\tau \in \mathcal{D}_\sigma(0,100) \setminus T(\boldsymbol{\Phi})$, then τ leaves at least $\mathcal{U}_{\boldsymbol{\Phi}}(\tau) \geq \rho n/10$ clauses unsatisfied. More precisely, given $\mathcal{U}_{\boldsymbol{\Phi}}(\tau)$ each unsatisfied clause consists of k independent random literals that are unsatisfied under τ. Since $\mathcal{D}_\sigma(0,100)$, for any one of the $k\mathcal{U}_{\boldsymbol{\Phi}}(\tau)$ underlying variables the probability of belonging to $\Delta(\sigma,\tau)$ equals $\Delta(\sigma,\tau)/n \leq 100\kappa$. Therefore, Lemma 1.3 shows that

$$\mathbb{P}\left[\mathcal{P}_{\boldsymbol{\Phi}}(\sigma,\tau)\right] \leq \mathbb{P}\left[\mathrm{Bin}(k|\mathcal{U}_{\boldsymbol{\Phi}}(\tau)|, \Delta(\sigma,\tau)/n) > k\mathcal{U}_{\boldsymbol{\Phi}}(\tau)/10\right] \leq \exp(-k\rho n/10). \tag{4.10}$$

Summing (4.10) on $\sigma \in \{0,1\}^n$ and $\tau \in \mathcal{D}_\sigma(0,100)$ and using our assumption on ρ, we get

$$\mathbb{E}\left[\mathcal{P}\right] \leq \sum_{\sigma,\tau} \mathbb{P}\left[\mathcal{P}_{\boldsymbol{\Phi}}(\sigma,\tau)\right] \leq 4^n \exp(-k\rho n/10) \leq 2^{-n}$$

Thus, the assertion follows from Markov's inequality. □

Finally, Proposition 2.2 follows directly from Lemma 4.2 to 4.4.

References

[1] D. Achlioptas: Lower bounds for random 3-SAT via differential equations. Theoretical Computer Science **265** (2001) 159–185.
[2] D. Achlioptas, P. Beam, M. Molloy: Exponential bounds for DPLL below the satisfiability threshold. Proc. SODA (2004).
[3] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.
[4] D. Achlioptas, A. Coja-Oghlan, F. Ricci-Tersenghi: On the solution-space geometry of random constraint satisfaction problems. Random Structures and Algorithms **38** (2011) 251–268.
[5] D. Achlioptas, C. Moore: Random k-SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
[6] D. Achlioptas, Y. Peres: The threshold for random k-SAT is $2^k \ln 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.
[7] M. Alekhnovich, E. Ben-Sasson: Linear upper bounds for random walk on small density random 3-CNFs. SIAM Journal on Computing **36** (2006) 1248–1263.
[8] M. Alekhnovich, E. Ben-Sasson: Analysis of the random walk algorithm on random 3-CNFs. unpublished (2002).
[9] A. Braunstein, L. Dall-Asta, G. Semerjian, L. Zdeborová: The large deviations of the whitening process in random constraint satisfaction problems. Journal of Statitstical Mechanics: Theory and Experiment **5** (2016) 053401.
[10] A. Broder, A. Frieze, E. Upfal: On the satisfiability and maximum satisfiability of random 3-CNF formulas. Proc. 4th SODA (1993) 322–330.
[11] M.-T. Chao, J. Franco: Probabilistic analysis of a generalization of the unit-clause literal selection heuristic for the k-satisfiability problem. Inform. Sci. **51** (1990) 289–314.
[12] P. Cheeseman, B. Kanefsky, W. Taylor: Where the *really* hard problems are. Proc. IJCAI (1991) 331–337.
[13] V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.
[14] A. Coja-Oghlan: A better algorithm for random k-SAT. SIAM J. Computing **39** (2010) 2823–2864.
[15] A. Coja-Oghlan: On belief propagation guided decimation for random k-SAT. Proc. 22nd SODA (2011) 957–966.
[16] A. Coja-Oghlan, A. Frieze: Analysing Walksat on random formulas. SIAM Journal on Computing **43** (2014) 1456–1485
[17] A. Coja-Oghlan, K. Panagiotou: The asymptotic k-SAT threshold. Advances in Mathematics **288** (2016) 985–1068.
[18] E. Dantsin, A. Wolpert: An improved upper bound for SAT. Proc. 8th SAT (2005) 400–407.
[19] J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k. Proc. 47th STOC (2015) 59–68.
[20] A. Frieze, S. Suen: Analysis of two simple heuristics on a random instance of k-SAT. Journal of Algorithms **20** (1996) 312–355.
[21] D. Gamarnik, M. Sudan: Limits of local algorithms over sparse random graphs. Proc. of 5th ICTS (2014) 369–376.
[22] D. Gamarnik, M. Sudan: Performance of Survey Propagation guided decimation algorithm for the random NAE-K-SAT problem. arXiv 1402.0052v2 (2014).
[23] M. Hajiaghayi, G. Sorkin: The satisfiability threshold of random 3-SAT is at least 3.52. IBM Research Report RC22942 (2003).
[24] T. Hertli: 3-SAT Faster and Simpler - Unique-SAT Bounds for PPSZ Hold in General. SIAM J. Comput. **43** (2014) 718–729.
[25] T. Hertli, R. Moser, D. Scheder: Improving PPSZ for 3-SAT using critical variables. Proc. 28th STACS (2011) 237–248.
[26] S. Hetterich: Analysing Survey Propagation Guided Decimation on Random Formulas. Proc. of 43rd ICALP (2016) in press.
[27] T. Hofmeister, U. Schöning, R. Schuler, O. Watanabe: A Probabilistic 3-SAT Algorithm Further Improved. Proc. 19th STACS (2002) 192–202.

[28] K. Iwama, S. Tamaki: Improved Upper Bounds for 3-SAT. Proc. 15th SODA (2004) 328–328.
[29] S. Janson, T. Łuczak, A. Ruciński: Random Graphs, Wiley 2000.
[30] A. Kaporis, L. Kirousis, E. Lalas: The probabilistic analysis of a greedy satisfiability algorithm. Random Structures and Algorithms **28** (2006) 444–480.
[31] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.
[32] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.
[33] D. Mitchell, B. Selman, H. Levesque: Hard and easy distribution of SAT problems. Proc. 10th AAAI (1992) 459–465.
[34] M. Molloy: Cores in random hypergraphs and Boolean formulas. Random Struct. Algorithms **27** (2005) 124–135.
[35] C. H. Papadimitriou: On selecting a satisfying truth assignment. Proc. 32nd FOCS (1991) 163–169.
[36] R. Paturi, P. Pudlák, M. Saks, F. Zane: An Improved Exponential-time Algorithm for k-SAT. J. ACM **52** (2005) 337–364.
[37] M. Rahman, B. Virag: Local algorithms for independent sets are half-optimal. Annals of Probability, in press.
[38] U. Schöning: A probabilistic algorithm for *k*-SAT and constraint satisfaction problems. Proc. 40th FOCS (1999) 410–414.
[39] B. Selman, H. Kautz, B. Cohen: Local search strategies for satisfiability testing. In David S. Johnson, Michael A. Trick (eds.): Cliques, coloring, and satisfiability: second DIMACS implementation challenge, October 11-13, 1993. DIMACS Series in Discrete Mathematics and Theoretical Computer Science **26** (1996).
[40] G. Semerjian, R. Monasson: A study of pure random walk on random satisfiability problems with "physical" methods. Proc. 6th SAT (2003) 120–134.

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

AMIR HAQSHENAS, amirhaqshenas@googlemail.com, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

SAMUEL HETTERICH, hetterich@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.